arXiv:1505.07192v1 [cs.CV] 27 May 2015

Inner and Inter Label Propagation: Salient Object Detection in the Wild

Hongyang Li, *Student Member, IEEE,* Huchuan Lu, *Senior Member, IEEE,*
Zhe Lin, *Member, IEEE,* Xiaohui Shen, *Member, IEEE,* and Brian Price, *Member, IEEE*

***Abstract*— In this paper, we propose a novel label propagation based method for saliency detection. A key observation is that saliency in an image can be estimated by propagating the labels extracted from the most certain background and object regions. For most natural images, some boundary superpixels serve as the background labels and the saliency of other superpixels are determined by ranking their similarities to the boundary labels based on an inner propagation scheme. For images of complex scenes, we further deploy a 3-cue-center-biased objectness measure to pick out and propagate foreground labels. A co-transduction algorithm is devised to fuse both boundary and objectness labels based on an inter propagation scheme. The compactness criterion decides whether the incorporation of objectness labels is necessary, thus greatly enhancing computational efficiency. Results on five benchmark datasets with pixel-wise accurate annotations show that the proposed method achieves superior performance compared with the newest state-of-the-arts in terms of different evaluation metrics.**

***Index Terms*— Label Propagation, Saliency Detection.**

I. Introduction

HUMANS have the capability to quickly prioritize external visual stimuli and localize their most interested regions in a scene [1]. In recent years, visual attention has become an important research problem in both neuroscience and computer vision. One branch focuses on eye fixation prediction [2], [3], [4] to investigate the mechanism of human visual systems whereas the other trend concentrates on salient object detection [5], [6], [7] to accurately identify a region of interest. Saliency detection has served as a pre-processing procedure for many vision tasks, such as collages [8], image compression [9], stylized rendering [10], object recognition [11], visual tracking [12], image retargeting[13], etc.

In this work, we focus on the salient object detection. Recently, many low-level features directly extracted from images have been explored. It has been verified that color contrast is a primary cue for satisfying results [5], [10]. Other representations based on the low-level features try to exploit the intrinsic textural difference between the foreground and background, including focusness [1], textual distinctiveness [14], and structure descriptor [15]. They perform well in many cases, but can still struggle in complex images. Instead, we observe that the primitive appearance information alone is good enough to reflect the textural difference from the boundaries of superpixels.

Due to the shortcomings of low-level features, many algorithms have turned to incorporating higher-level features [16], [17], [18], [19]. One type of higher-level representations that can be employed is the notion of objectness [20], or how likely a given region is an object. For example, Jiang *et al.* [1] compute a saliency measure by combining the objectness values of many overlapping windows. However, using the objectness measure directly to compute saliency may produce unsatisfying results in complex scenes when the objectness score fails to predict true salient object regions [21], [22]. A better way to employ high-level objectness is to consider the scores as hints of the foreground.

To this end, we put forward a unified approach to incorporate low-level features and the objectness measure for saliency detection via label propagation. Since the border regions of the image are good indicators to distinguish salient objects from the background [23], [24], we observe that the boundary cues can be used to estimate the appearance of the background while the objectness cues focus on the characteristics of the salient object. Therefore, a refined co-transduction [25] based method, namely *label propagation saliency* (LPS), is proposed. In this framework, the most certain boundary and object regions are able to propagate saliency information in order to best leverage their complementary influence. As the boundary cue can be quite effective in some cases and the objectness measure requires additional computation, a compactness criterion is further devised to determine whether the results propagated by boundary labels are sufficient.

Fig.1 shows the pipeline of our method. First, we extract the affinity matrix and choose some border nodes as labels to represent the background (Sec.III-A). The *inner* propagation is implemented to obtain the regional maps (Sec.III-B). Second, a compactness criterion is introduced to evaluate whether these maps need a further refinement (Sec.III-D). Third, the *inter* propagation incorporates objectness labels via a co-transduction algorithm to regenerate maps for images that fail to work in the inner stage (Sec.III-C, III-D). Fourth, all maps are updated at a pixel level to achieve coherency of the saliency assignment (Sec.III-E). The contributions of our work include:

This work was supported by the Natural Science Foundation of China #61472060 and the Fundamental Research Funds for Central Universities under Grant DUT14YQ101. We thank Lijun Wang for helpful discussions along the way; Na Tong and Wei Liu for careful proofreading and reviewers for fruitful and detailed comments.

H. Li is affiliated with the Department of Electronic Engineering, The Chinese University of Hong Kong. This research project was done when he was a final-year undergraduate and served as a research assistant in the computer vision group at Dalian University of Technology. E-mail: hongyangli2020@gmail.com.

H. Lu is with the School of Information and Communication Engineering, Dalian University of Technology. Email: lhchuan@dlut.edu.cn.

Z. Lin, X. Shen and B. Price are with Adobe Research, San Jose, USA. Emails: {zlin, xshen, bprice}@adobe.com

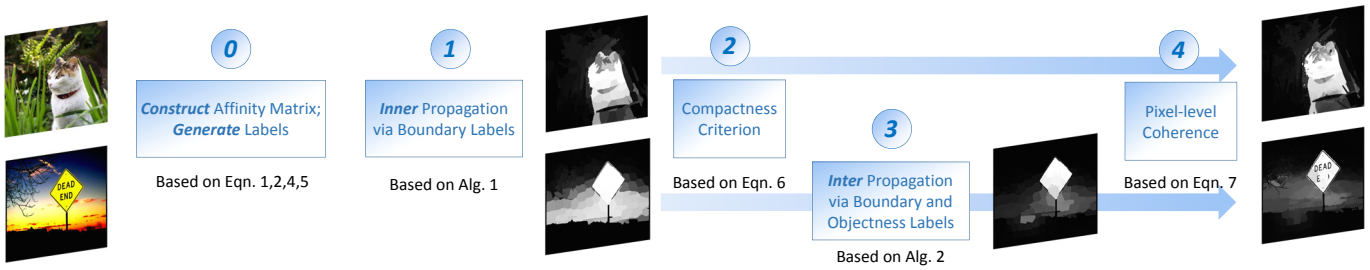

Fig. 1. Pipeline of the label propagation saliency algorithm. First, we construct the normalised affinity matrix from the superpixels and generate boundary and objectness label sets, respectively; then the inner propagation is conducted to have initial saliency maps; third, the compactness criterion chooses those who need a further refinement by the inter propagation scheme; finally, all maps are enhanced via a pixel-level saliency coherence.

1) A simple and efficient label propagation algorithm via boundary labels for most natural images based on the reconstructed affinity matrix;
2) A novel co-transduction framework to incorporate foreground labels obtained from the objectness measure with boundary cues for complex images;
3) A compactness selection mechanism to decide whether the initial maps need an update, thus facilitating the computational efficiency.

The experimental results show that the proposed method achieves superior performance in various evaluation metrics against other 27 state-of-the-arts on five image benchmarks. Finally, the results and code are shared for research purposes[1].

II. Related Work

Saliency estimation methods can be explored from different perspectives. Basically, most works employ a bottom-up approach via low level features while a few incorporate a top-down solution driven by specific tasks. Early researches address saliency detection via biologically inspired models, such as Gaussian pyramids [2], fuzzy growing [26], graph-based activation [27]. Other studies employ frequency domain methods [28], [29], [30] to determine saliency according to the spectrum of the image's Fourier transform. However, the results of these methods exhibit undesirable blurriness and tend to highlight object boundaries rather than its entire area.

Recently, the saliency detection community has witnessed a blossom of high accuracy results under distinctive frameworks [31]. Learning methods [19], [17], [13] integrate both low and high level features to compute saliency based on parameters trained from sample images. Although learning mechanisms perform well in proposing bounding boxes, they suffer in salient object detection due to the complex scenes of the background. Shen *et.al* [16] introduce high-level priors to form high-dimensional representations of the image and construct saliency in a low rank framework. Despite the complicated configuration, the resultant maps have unsatisfying saliency assignment near the salient object.

Faced with the above issues and considering the limited knowledge of structural description mentioned in Sec.I, we try to extract features in a simple and effective way. Jiang *et al.* [24] introduce an absorbing Markov chain method where the appearance divergence and spatial distribution between salient objects and the background are considered. Cheng *et al.* [10] formulate a regional contrast based saliency algorithm which simultaneously evaluates global and local contrast differences. Inspired by these works, we construct an affinity matrix based on the color feature of superpixels with two adjustments to involve spatial relations.

A novel label propagation method is proposed in [32] to rank the similarity of data points to the query labels for shape retrieval. We apply and refine the theory to make full use of the background and foreground superpixels, which has been rarely studied in saliency detection. Distinct from the work of Yang *et al.* [23] where a manifold ranking algorithm assigns saliency based on priors of all boundary nodes, in this work, (a) we only take some boundary nodes to eliminate salient regions that appear at the image border; (b) both boundary and foreground nodes are selected as complementary labels in a co-transduction framework to fully distinguish salient areas from the background; and (c) the revised label propagation algorithm has zero parameter whereas in [23] the sensitive α has a vital effect on results in different datasets.

III. The Label Propagation Algorithm

We first introduce the construction of the affinity matrix in Sec.III-A, which is of vital importance during the label propagation. Then the inner propagation via boundary labels is proposed in Sec.III-B. An objectness measure is utilised to locate foreground labels in Sec.III-C. Sec.III-D illustrates the co-transduction algorithm which takes into consideration both boundary and objectness cues and the compactness criterion to classify initial maps generated from the inner propagation. Finally, we refine the regional maps on pixel level to achieve saliency coherency in Sec.III-E.

A. Affinity Matrix Construction

We first construct an affinity matrix among superpixels to be used in the propagation algorithm. L_0 gradient minimization [33] is implemented to obtain a soft abstraction layer while keeping vital details of the image. Superpixels are generated to segment the smoothed image into N regions by the SLIC algorithm [34], where regions at the image border form a set

[1] https://github.com/hli2020/lps_tip15

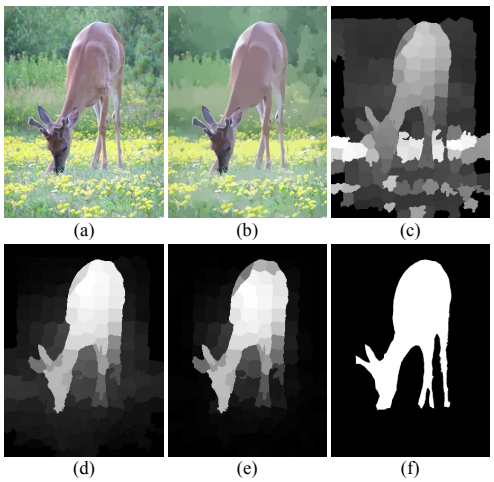

Fig. 2. Effects of affinity construction. (a) Input image; (b) L_0 smoothing; (c) Full connection; (d) No geodesic constraint; (e) Geodesic constraint; (f) Ground truth.

of boundary nodes, denoted as B. In this work, we refer the superpixel as a node or a region.

The similarity of two nodes is measured by a defined distance of the mean features in each region. Based on the intuition that neighboring regions are likely to share similar appearances and that remote ones do not bother to have similar saliency values even if the appearance of them are highly identical, we define the affinity entry w_{ij} of superpixel i to a certain node j as:

$$w_{ij} = \begin{cases} \exp(-\frac{D(\mathbf{f}_i,\mathbf{f}_j)}{\sigma^2}) & j \in \mathcal{N}(i) \text{ or } i,j \in B \\ 0 & i = j \text{ or otherwise} \end{cases} \quad (1)$$

where $\mathbf{f}_i, \mathbf{f}_j$ denote the mean feature vectors of pixels inside node i, j respectively, σ is a tuning parameter to control strength of the similarity, $\mathcal{N}(i)$ indicates the set of the direct neighboring nodes of superpixel i, as well as the direct neighbors of those neighboring nodes. Therefore, we have an affinity matrix $\mathbf{W} = [w_{ij}]_{N\times N}$ to indicate the similarity between any pair of superpixels, a degree matrix $\mathbf{D} = diag\{d_1, \ldots, d_N\}$ where $d_i = \sum_j w_{ij}$ to sum the total entries of each node to other nodes, and a row-normalized affinity matrix:

$$\mathbf{A} = \mathbf{D}^{-1} \cdot \mathbf{W} \quad (2)$$

to be finally adopted.

Different from the common practice of a fully connected network among superpixels [10], [22], there are two adaptions to construct the affinity entry in Eqn.1. First, a conception of k-layer neighborhood (here $k = 2$) in graph theory is introduced. The enlarged neighbors of the region enforce a spatial relationship that salient object tends to be clustered rather than to be scattered. Second, we adopt a geodesic constraint mechanism [24], [23] to further enhance the relationship among boundary nodes, *i.e.*, any two superpixels in B are connected. Since the boundary nodes serve as propagation labels, a strong connection among them could better distinguish the background from the salient object.

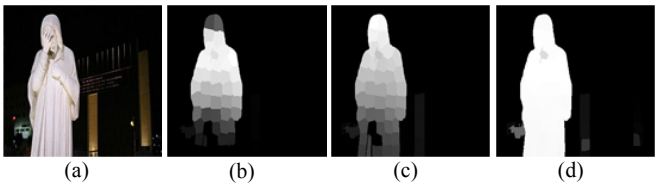

Fig. 3. Salient objects at image border. (a) Input image; (b) Use all boundary nodes; (c) No geodesic constraint; (d) Selected boundary nodes.

The effects of affinity construction are illustrated in Fig.2. We note that under a fully connected scheme in Fig.2(c), yellow flowers in the background are salient due to the mere consideration of color and ignorance of spatial distance. Without a geodesic constraint scheme, the saliency map in Fig.2(d) has vast background areas with low saliency assignment which leads to a low precision at a high recall in the precision-recall curve.

B. Inner Propagation via Boundary Labels

Given an affinity matrix, we endeavor to propagate the information of the background labels to estimate saliency measure of other superpixels. A shape similarity method that exploits the intrinsic relation between labelled and unlabelled objects is proposed in [32] to tackle the image retrieval problem via label propagation. Given a dataset $R = \{r_1, \ldots, r_l, r_{l+1}, \ldots, r_N\} \in \mathbb{R}^{D\times N}$, where the former l regions serve as query labels and D denotes the feature dimension, we seek out a function $\mathbf{V} = [V(r_1), \ldots, V(r_N)]^T$ such that $\mathbf{V} : R \to [0,1] \in \mathbb{R}^{N\times 1}$ indicates the possibility of how similar each data point is to the labels. The similarity measure $V(r_i)$ satisfies

$$V_{t+1}(r_i) = \sum_{j=1}^{N} a_{ij} V_t(r_j) \quad (3)$$

where a_{ij} is the affinity entry defined in Eqn.2 and t is the recursion step.

The similarity measure of query labels is fixed to be 1 during the recursive process and the initial measure of unlabelled objects is set to be 0. For a given region, the similarity $V(r_i)$ is learned iteratively via propagation of the similarity measures of its neighbors $V(r_j)$ such that a region's final similarity to the labels is effectively influenced by the features of its surroundings. In other words, the new similarity will be large iff all points r_j that resemble r_i are also quite similar to query labels. Fig.4 shows a simple example on how Eqn.3 plays a vital role in the saliency propagation process.

Specifically, we choose CIE LAB color as the input feature because distance in LAB space matches human perception well. A color-based affinity matrix $\mathbf{A}_c$ with a controlling parameter σ_c is constructed according to Eqn.1, where the feature distance $D(\mathbf{f}_i, \mathbf{f}_j) = \|\mathbf{c}_i - \mathbf{c}_j\|_2$. The boundary nodes are employed as the query labels simply because regions near the image border are less likely to be salient. However, as is shown in Fig.3(b), in some cases, the salient object appears at the border and the saliency measure is doomed to be 0 if the salient region is chosen to be the background labels. Consequently, we compute the color distinctiveness of each boundary node from other border regions according to Eqn.1,

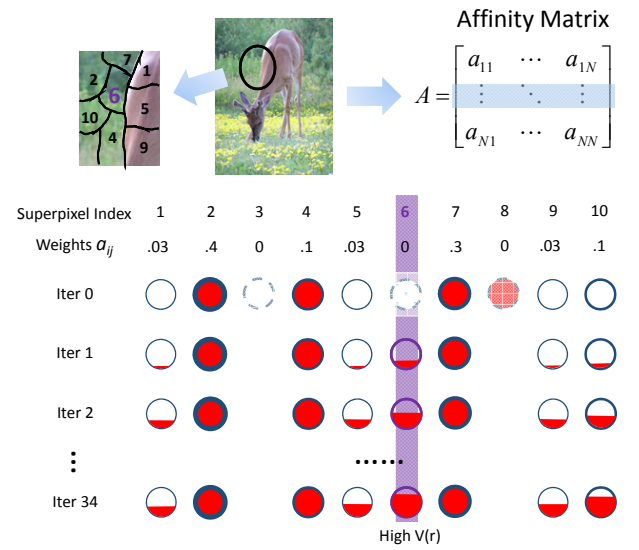

Fig. 4. A toy example to illustrate how the inner propagation algorithm (Alg.1), or Eqn.3, works. For simplicity, we investigate one superpixel region (#6) and see how its value $V(r)$ changes during each iteration. Assuming we have 10 regions in total and the 6-*th* row of the normalised affinity matrix (weights a_{ij}) shows the similarity between the considered region and other regions. The dash-outline regions ($\#3, 6, 8$) are not neighbours of region 6 and thus not considered in the propagation. The outline weight of each circle indicates the affinity weight, *i.e.*, the thicker it looks, the bigger a_{ij} is. The red area inside each circle denotes the value of $V(r)$, since we assume region 2,4,7 are background labels, they have red colour fully filled within their circles in each iteration.

drop the top 30% with high color difference empirically, and thus create the set of selected boundary labels B'. We can also observe from Fig.3(c) that the geodesic constraint with selected boundary labels facilitates saliency accuracy in such scenarios by strengthening the connection among boundary regions.

Alg.1 summarizes the inner label propagation via boundary nodes. The convergence of the similarity measure $\mathbf{V}$ is ensured by checking whether its average variance in the last 50 iterations (*i.e.*, $const = 49$) is below a threshold. $\texttt{sp2map}(\cdot)$ means mapping the saliency measures of N regions into an image-size map. Note that such a propagation framework is similar to that in [35]. However, we find the ranking results obtained from a closed-form solution less encouraging than those of ours due to different constructions of the affinity matrix.

In most cases, the inner propagation with help of the boundary labels works well whereas in some complex scenes, as is shown in Fig.5(d), depending on the boundary prior alone might lead to high saliency assignment to the background regions. It naturally suggests us to use some foreground prior to improve the results further.

C. Objectness Labels as Foreground Prior

Alexe *et al.* [20] propose a novel method based on low-level cues to compute an objectness score for any given image window, which indicates the likelihood of the window containing an object. Several useful priors are exploited and combined in a Bayesian framework, including multi-scale saliency (MS), color-contrast (CC), edge density (ED), superpixel straddling

Algorithm 1 Inner Label Propagation via Boundary Nodes

Input:
The $N \times N$ row-wise normalized color affinity matrix $\mathbf{A}_c$. The set of selected boundary labels B' and the set of unlabelled nodes $U = \{R \backslash B'\}$.

$t = 0$
Initialize, set $V_t(r_i) = 1$ for $r_i \in B'$ and $V_t(r_i) = 0$ for $r_i \in U$
while $check > thres$ **do**
  for $r_i \in U$ **do**
    $V_{t+1}(r_i) = \sum_{j=1}^{N} a_{ij} V_t(r_j)$
  end for
  $t = t + 1$
  $check = \texttt{var}(\mathbf{V}_t, \mathbf{V}_{t-1}, \ldots, \mathbf{V}_{t-const})$
end while
$\mathbf{S}^B = \texttt{ones}(N) - \texttt{normalize}(\mathbf{V}_t)$
$S^B(r_i) = \texttt{sp2map}(\mathbf{S}^B)$

Output:
The regional map $S^B(r_i)$ from background labels.

(SS) and location plus size (LS). The results show high performance on the PASCAL VOC 07 dataset.

- MS, proposed by [28], measures the uniqueness of objects according to the spectral residual of the image's FFT.
- CC, similar as in [5], considers the distinct appearance of objects via a center-surround histogram of color distribution.
- ED and SS capture the closed boundary of objects. The former computes the density of edges near window borders while the latter calculates how intact superpixels are inside a window.
- LS exploits the likeliness of a window to cover an object based on its size and location using kernel density estimation.

In practice, we find the first three cues more important while the last two more trivial. Big and homogeneous superpixels are generated by [36] in [20] whereas small, tiny, and compact superpixels are created by the SLIC algorithm, making SS incompatible in our work. Furthermore, LS measures the size and location of windows without taking into consideration the intrinsic features of images and often dominates the final integrated objectness score due to different image benchmarks. To this end, we only utilize MS, CC and ED since cues are combined independently in a naive Bayes model. The rest of the parameters in the objectness measure are set to be default as in [20].

Let P_m be a probability score of the m-th sampling window, the pixel-level objectness map $\mathcal{O}(p)$ is obtained through overlapping scores multiplied by the Gaussian smoothing kernel of all sampling windows:

$$\mathcal{O}(p) = \sum_{m=1}^{M} P_m \cdot \exp\left[-\left(\frac{(x_p - x_m^c)^2}{2\sigma_x^2} + \frac{(y_p - y_m^c)^2}{2\sigma_y^2}\right)\right] \quad (4)$$

where $M = 1000$ is the number of sampling windows, x_p, y_p, x_m^c, y_m^c denote the coordinates of pixel p and the center coordinates of window m respectively. We set $\sigma_x = 0.25W$

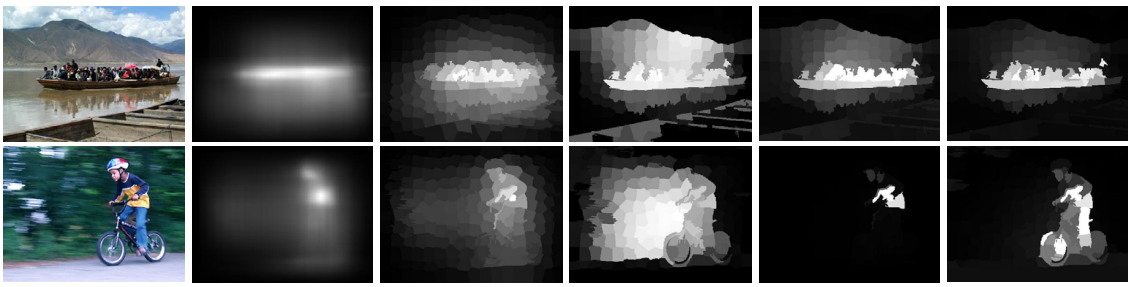

Fig. 5. Objectness integration. (a) Input image; (b) Pixel-level objectness map; (c) Region-level objectness map; (d) Inner boundary propagation; (e) Inner objectness propagation; (f) Inter propagation via boundary and objectness labels.

and $\sigma_y = 0.25H$, where W is the width and H the height of an image. The region-level objectness map $\mathcal{O}(r_i)$ is the average of pixels' objectness values within a region:

$$\mathcal{O}(r_i) = \frac{1}{n_i} \sum_{p \in r_i} \mathcal{O}(p) \tag{5}$$

where n_i indicates the number of pixels in region r_i.

The integration of objectness labels is illustrated in Fig.5. By introducing only three cues of the objectness measure and a Gaussian kernel refinement, the pixel-level map in Fig.5(b) can better capture and highlight the focus of a salient object. The region-level objectness map in Fig.5(c) is obtained similarly as one of the three saliency maps in [1]. A simple average of pixels' scores within a region leads to mid-value saliency in vast background areas since the pixel-level map from which the region-level map is generated is ambiguous around the salient object in the first place.

Based on the fact that high values of region-level objectness score calculated by Eqn.5 can better indicate foreground areas, the set of objectness labels O is created from superpixels whose region-level objectness $\mathcal{O}(r_i)$ is no less than the objectness criterion γ_1. Fig.5(e) displays the saliency maps by the inner label propagation via objectness labels alone. We observe that under the objectness mechanism, the top image effectively inhibits high values of the background saliency while the bottom image only detects the kid's orange shirt due to a limited number of label hints from set O. This indicates that a complementary combination of the boundary and objectness labels could be a better choice.

D. Inter Propagation via Co-transduction

Recently, Bai *et al.* [25] propose a similarity ranking algorithm by fusing different affinity measures for robust shape retrieval under a semi-supervised learning framework. Inspired by such an idea, we devise a new co-transduction algorithm for saliency detection, which uses one label set to pull out confident data and add additional labels as new hints to the other label set. The inter label propagation algorithm is summarized in Alg.2. Besides different application areas, our algorithm differentiates from the original work [25] in the following three ways:

First, instead of fusing two different similarity matrices, we construct the same matrix $\mathbf{A}_c$ for both label sets (through line 7 to 8). Fusing two affinity matrices is investigated and an orientation-magnitude (OM) descriptor [15] is extracted to capture the structural characteristic of images. We compute a structure-based affinity matrix $\mathbf{A}_s$ according to Eqn.1, where $D(\mathbf{f}_i, \mathbf{f}_j) = \chi^2(h_{OM}(i), h_{OM}(j))$ and $\sigma_s^2 = 0.1$. As shown in Fig.6(a)-(c), the saliency map using one color affinity matrix outperforms that of using two matrices. The information from structure description seems to be redundant since the color affinity matrix $\mathbf{A}_c$ already includes knowledge of textual

Algorithm 2 Inter Label Propagation via Boundary and Objectness Nodes

Input:
The $N \times N$ row-wise normalized color affinity matrix $\mathbf{A}_c$. The set of selected boundary labels B' and the set of objectness labels O.

1: $t = 0$
2: Initialise, $\mathbf{V}_t^B = \mathbf{0}, \mathbf{V}_t^O = \mathbf{0}$
3: **while** $check^B, check^O > thres$ **do**
4: Set $V_t^B(r_i) = 1$ for $r_i \in B'$, $V_t^O(r_i) = 1$ for $r_i \in O$
5: Create unlabelled sets U_1 and U_2 such that $U_1 = \{R \backslash B'\}, U_2 = \{R \backslash O\}$
6: **for** $r_i \in U_1, r_i \in U_2$ **do**
7: $V_{t+1}^B(r_i) = \sum_{j=1}^N a_{ij} V_t^B(r_j)$
8: $V_{t+1}^O(r_i) = \sum_{j=1}^N a_{ij} V_t^O(r_j)$
9: **end for**
10: $t = t + 1$
11: $check^B = \mathtt{var}(\mathbf{V}_t^B, \ldots, \mathbf{V}_{t-const}^B)$
12: $check^O = \mathtt{var}(\mathbf{V}_t^O, \ldots, \mathbf{V}_{t-const}^O)$
13: $temp1 = \mathtt{sort}(\mathbf{V}_t^B, \text{'ascend'})$
14: $temp2 = \mathtt{sort}(\mathbf{V}_t^O, \text{'ascend'})$
15: $L^B = temp1(1 : p_1)$, $L^O = temp2(1 : p_2)$
16: $B' = B' \cap L^O$, $O = O \cap L^B$
17: **end while**
18: $\mathbf{S}^B = \mathtt{ones}(N) - \mathtt{normalize}(\mathbf{V}_t^B)$
19: $\mathbf{S}^O = \mathtt{normalize}(\mathbf{V}_t^O)$,
20: $\mathbf{S}^C = \mathtt{normalize}(\alpha \mathbf{S}^B + \beta \mathbf{S}^O)$
21: $S^C(r_i) = \mathtt{sp2map}(\mathbf{S}^C)$

Output:
The combined regional saliency map $S^C(r_i)$.

distinctiveness at the borders of each region.

Second, we emphasize more on the difference between boundary and objectness labels during the propagation whereas the same p labels are switched within each label set in [25]. During each iteration in Alg.2 (through line 11 to 16), p_1 superpixels which are most different from the boundary labels are picked out and added to the objectness set and the update of the boundary set is similarly achieved with a different superpixel number p_2. We set p_1, p_2 to be $p_1 \ll p_2$ because the background regions often significantly outnumber the foreground ones.

Third, we observe that the ranking values in the first few recursions to be highly noisy and inaccurate. Therefore, unlike the practice of [25] that averages all the similarity measures in each iteration, the final saliency measure is computed as a linear combination of the resultant $\mathbf{S}^B$ and $\mathbf{S}^O$ in the last iteration from boundary and objectness labels, respectively (through line 18 to 21).

From Fig.5(d)-(f) we can see that such a co-transduction algorithm outperforms the inner label propagation via boundary or objectness nodes alone. The failure images in the inner boundary propagation are often cases where there are vast areas with high-value saliency assignment to regions around the salient object. The reason ascends from the resemblance of appearance between salient and non-salient nodes, as well as the spatial discontinuity from the boundary labels to the center regions which prevent labels from being propagated to the background regions around the image center. The inter propagation algorithm strengthens the connection of salient regions by employing objectness labels and distinguishes the foreground better from the background by enlarging the set of boundary labels from objectness cues, thus best leveraging the complimentary information of both label sets. Some may argue that the graph cuts optimization method also performs well in saliency detection [37], but the co-transduction algorithm is designed to obtain continuous saliency assignment while the former aims at solving binary MRF problems.

In some cases, as shown in Fig.6(d)-(f), the inner propagation via boundary labels alone has better saliency maps than a combination of boundary and objectness labels, which results from the slight disturbance of objectness measures near the salient object. To this end, we propose a compactness score to evaluate the quality of the regional saliency map $S^B(r_i)$ generated by Alg.1:

$$C(S) = \sum_{b=1}^{10} w(b) \cdot h^S(b) \tag{6}$$

where b denotes each quantisation of the resultant saliency map, $h^S(b)$ indicates a 10-bin histogram distribution of the map and $w(b)$ indicates the weight upon each bin. Based on the aforementioned characteristic of the failure saliency maps in the inner boundary propagation, we take a triangle form of the weight term, *i.e.*, $w(b) = \min(b, (11-b))$. Only the saliency maps with score lower than a compactness criterion γ_2 will be updated by the inter propagation via a co-transduction algorithm. Such a scheme not only ensures high quality of the saliency maps, but also improves the computational efficiency.

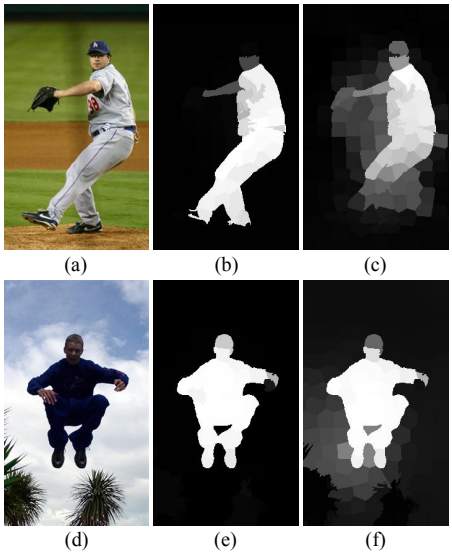

Fig. 6. Effects of co-transduction algorithm. (a) Input image; (b) Only color affinity matrix; (c) Both color and structure affinity matrix; (d) Input image; (e) Saliency map by Alg.1; (f) Saliency map by Alg.2.

E. Pixel-level Saliency Coherence

Finally, in order to eliminate the segmentation errors of the SLIC algorithm, we define the pixel-level saliency as a weighted linear combination of the regional saliency, $S^B(r_i)$ or $S^C(r_i)$, of its surrounding superpixels:

$$S(p) = \sum_{i=1}^{G} \exp\left(-(k_1\|\mathbf{c}_p - \mathbf{c}_i\| + k_2\|\mathbf{z}_p - \mathbf{z}_i\|)\right) S^{B/C}(r_i) \tag{7}$$

where $\mathbf{c}_p, \mathbf{c}_i, \mathbf{z}_p, \mathbf{z}_i$ are the color and coordinate vectors of a region or a pixel, G denotes the number of direct neighbors of region r_i, and S^B or S^C indicates the straightforward region-level result descending from Alg.1 or Alg.2. By choosing a Gaussian weight, we ensure the up-sampling process is both local and color sensitive. Here k_1 and k_2 are parameters controlling the sensitivity to color and position, where $k_1 = 0.2, k_2 = 0.01$ is found to work well in practice.

IV. Experimental Results

We evaluate the proposed method on five typical datasets. The first MSRA-1000, which is a subset of MSRA-5000, is a widely used dataset where almost every method has been tested by comparing to the accurate human-labelled masks provided in [29]. The second CCSD-1000 [38] contains more salient objects under complex scenes and some images come from the challenging Berkeley-300 dataset [39]. The third MSRA-5000 [5] includes a more comprehensive source of images with accurate masks recently released by [17]. The fourth THU-10,000 is the largest dataset in the saliency community so far where 10,000 images are randomly chosen from the MSRA database and the Internet with pixel-level labeling. The last PASCAL-S [40] ascends from the validation set of PASCAL VOC 2010 segmentation challenge. It contains 850 natural images where in most cases multiple objects of varying size, shape, color, etc., are surrounded by complex

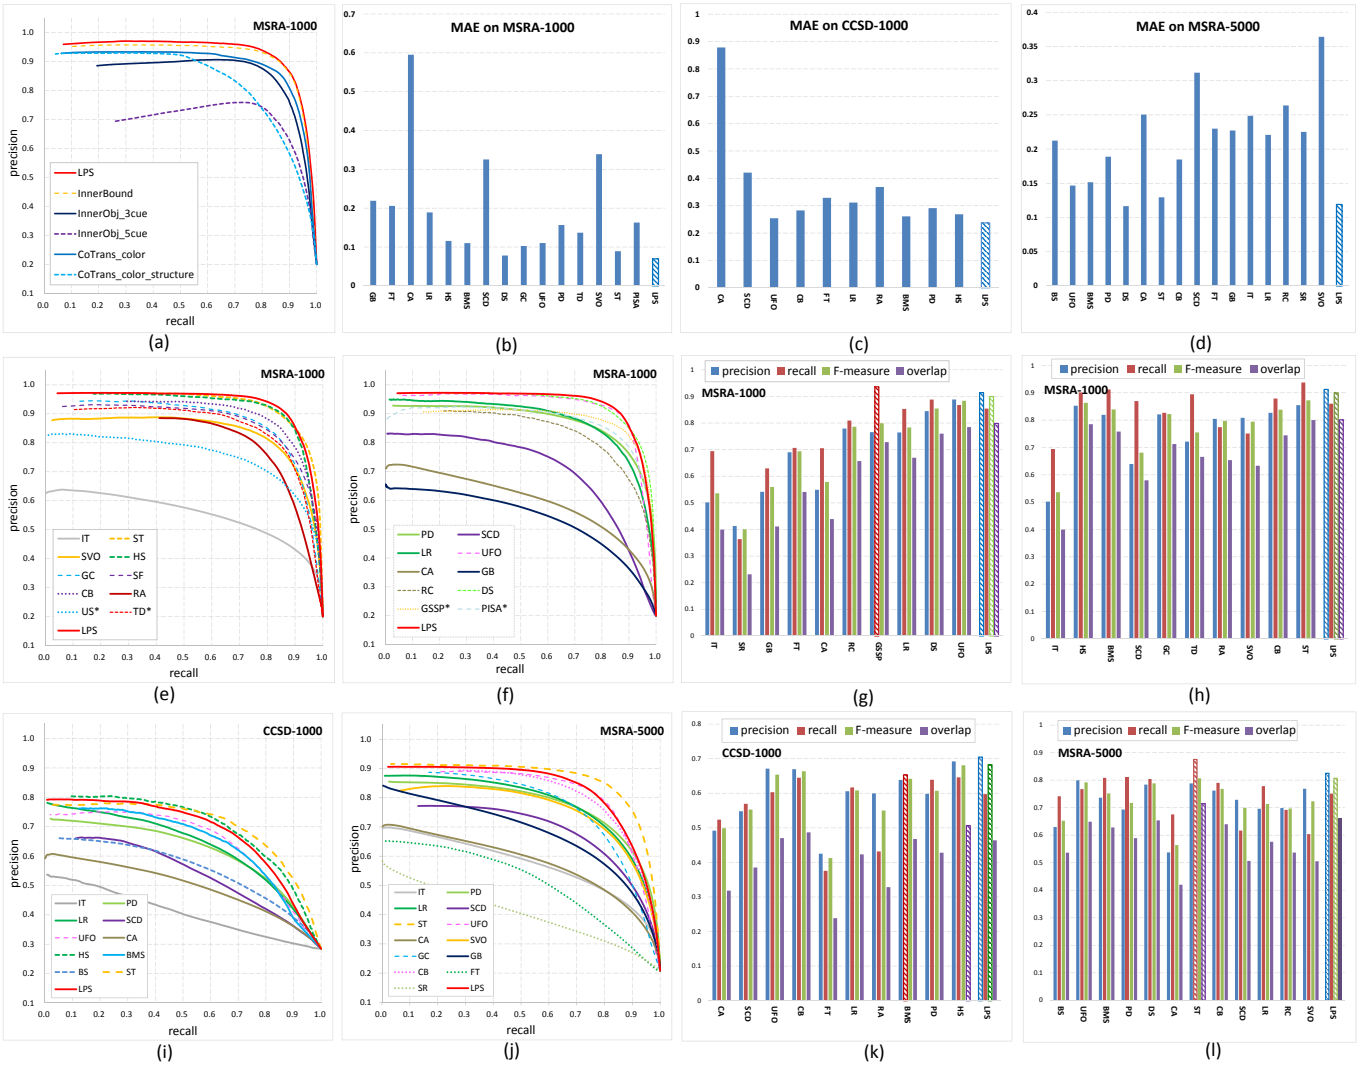

Fig. 7. Quantitative results. (a) Individual component analysis on MSRA-1000. Note that 'CoTrans_$*$' means implementing Alg.2 for every image; (b)-(d) *MAE* metric on MSRA-1000, CCSD-1000, MSRA-5000; (e)-(l) Performance comparison on MSRA-1000, CCSD-1000 and MSRA-5000 respectively. Bars with oblique lines denote the highest score in the corresponding metric. Methods followed by an asterisk (*) denote they are only compared in those datasets.

scenes. Unlike the traditional benchmarks, the PASCAL-S is believed to eliminate the dataset design bias.

The proposed LPS algorithm is compared with both the classic and newest state-of-the-arts: IT[2], GB[27], SR[28], LC[41], FT[29], CA[8], RA[42], CB[37], SVO[21], HC[10], BS[43], SF[44], LR[16], GSSP[39], MK[24], DS[45], GC[46], PD[47], MR[23], BMS[4], HS[38], US[19], UFO[1], TD[14], PISA[15], HPS[13], ST[48], SCD[49]. To evaluate these methods, we either use results provided by authors or run their implementations based on the available codes or softwares.

A. Parameters and Evaluation Metrics

1) Implementation Details: We set the control of color distance σ_c in Eqn.1 to be $\sigma_c^2 = 0.1$, the number of switching labels from background and objectness labels in Alg.2 to be $p_1 = 2, p_2 = 150$, respectively. The objectness criterion associated with Eqn.5 is chosen to be $\gamma_1 = 0.8$ and the compactness criterion with Eqn.6 is fixed at $\gamma_2 = 1.6$. Parameters are empirically selected (see Tab.II and Sec.IV-B7) and universally used for all images.

2) Fixed Threshold: In the first experiment we compare binary masks for every threshold in the range $[0, \ldots, 255]$ and calculate the precision and recall rate. Precision corresponds to the percentage of salient pixels correctly assigned, while recall corresponds to the fraction of detected salient pixels in relation to the number of salient pixels in ground truth maps.

3) Adaptive Threshold: In the second experiment we employ the saliency-map-dependent threshold proposed by [29] and define it as proportional to the mean saliency of a map:

$$T_a = \frac{k}{W \times H} \sum_{x=1}^{W} \sum_{y=1}^{H} S(x, y) \tag{8}$$

where k is typically chosen to be 1.5 [1]. Then a weighted harmonic mean measure between precision and recall, *i.e.*, F-measure, is introduced by

$$F_\beta = \frac{(1 + \beta^2) Precision \times Recall}{\beta^2 \times Precision + Recall} \tag{9}$$

where we set $\beta^2 = 0.3$ to emphasize precision [29]. As we can see later, one method cannot have in all the highest precision, recall and F-measure as the former two are mutually exclusive and the F-measure is a complementary metric to balance them.

Furthermore, the overlap rate R_o defined by the PASCAL VOC criterion (*i.e.,* intersection over union) is used to comprehensively leverage precision and recall under the adaptive-threshold framework.

4) Mean Absolute Error: In the third experiment we introduce the mean absolute error (MAE) between the continuous saliency map S and the binary mask of ground truth GT:

$$MAE = \frac{1}{W \times H} \sum_{x=1}^{W} \sum_{y=1}^{H} |S(x,y) - GT(x,y)|. \qquad (10)$$

The metric takes the true negative saliency assignments into account whereas the precision and recall favor the successfully assigned saliency to the salient pixels [46]. Moreover, the quality of the weighted continuous saliency maps may be of higher importance than the binary masks in some cases [44].

B. Quantitative Comparison

1) Individual Component Analysis: In order to demonstrate the effects of separate components and their combinations in our approach, we plot the precision-recall curves in Fig.7(a).

First, we see that the refined co-transduction algorithm (LPS) with a compactness selection mechanism outperforms the inner propagation via boundary labels or objectness ones alone. Second, the precision rate under the two-feature-matrices framework in the co-transduction (blue dashed line) goes down sharply at high recall, which indicates the structure descriptor cannot inhibit the background regions. Third, the take-all-cue scheme from [20] fails to achieve high precision especially at higher thresholds, which verifies our explanations in Sec.III-C to take only three cues. Note that in the inner objectness propagation, the precision at lower recall is even slightly worse because of the inaccurate objectness labels chosen from the non-salient regions.

2) Mean Absolute Error: Fig.7(b)-(d) shows the MAE metric of LPS and other methods on MSRA-1000, CCSD-1000 and MSRA-5000. Considering the recent and well-performed methods, such as DS13[45], GC13[46], BMS13[4], TD13[14], HS13[38], PISA13[15], PD13[47], LPS achieves the lowest error of 0.0695, 0.2369, 0.1191 on the corresponding datasets, which indicates the resultant maps have a high quality of highlighting salient objects while suppressing the background.

3) MSRA-1000: Fig.7(e)-(h) displays the P-R curves, F-measure and overlap rate on MSRA-1000 benchmark. On one hand, LPS achieves an average of 97% precision rate covering most ranges of the recall while models such as DS13[45], UFO13[1], HS13[38], ST[48], have similar performance competing ours and yet lower precision at specific ranges of recall; on the other hand, the highest precision, F-measure and overlap score of 0.91, 0.90, 0.80 is accomplished by LPS outperforming other 21 methods. Note that due to many false positive salient detections of GSSP[39], their model has the highest recall value in Fig.7(g); however, it is more important to have a high value of precision or F-measure in the saliency community.

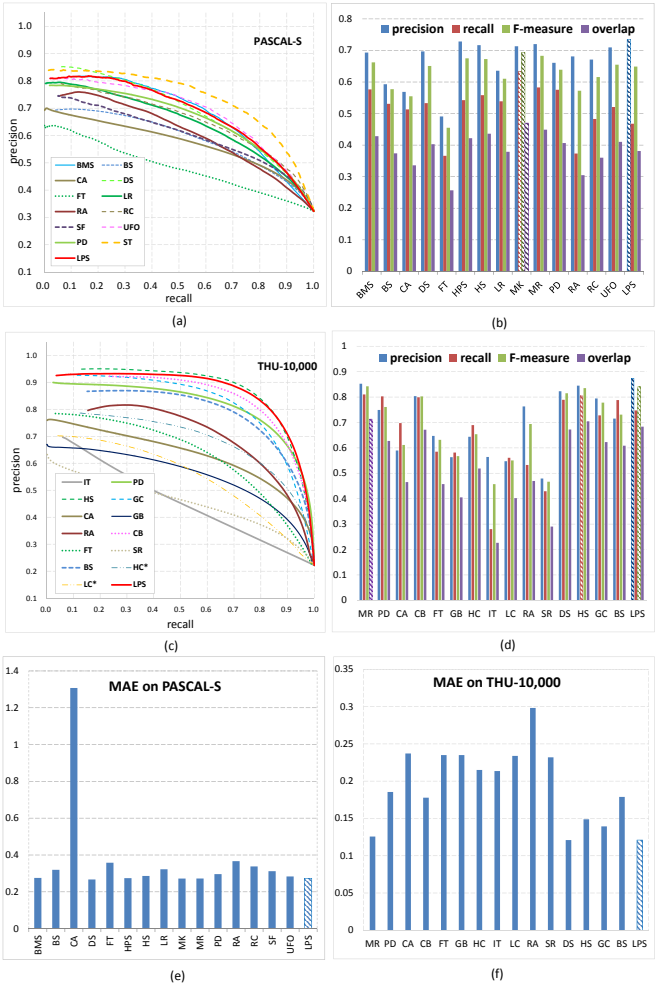

Fig. 8. Performance of the proposed algorithm compared with previous methods on the PASCAL-S (a),(b),(e) and THU-10,000 (c),(d),(f). Bars with oblique lines denote the highest score in the corresponding metric. Methods followed by an asterisk (*) denote they are only compared in those datasets.

4) CCSD-1000 and MSRA-5000: The last row of Fig.7 reports the performance comparison on these two datasets. For the CCSD benchmark, we observe that although HS13[38] achieves better precision curve and higher overlap score, LPS have the highest precision of 0.705, lowest MAE error and similar F-measure.

For the MSRA-5000, compared with most methods, LPS achieves the best curve performance spanning most ranges of recall as well as the highest precision and F-measure of 0.82, 0.81, respectively. We observe that the ST[48] model has a competitive high value of precision in the recall range from 0.7 to 1.0 (see Fig.7(j)), which means they have a strong capability to suppress the image background (even assigning small saliency value is not allowed). This advantage is probably attributed to the sentimental hierarchical analysis and the multi-scale scheme in their work. As for the adaptive threshold comparison, we have reached the highest precision and similar F-measure whereas ST keeps the highest overlap

TABLE I
EXECUTION TIME COMPARISON IN SECOND UNIT PER IMAGE ON THE MSRA-1000 DATASET. ALL CODES ARE DOWNLOADED FROM THE AUTHORS' WEBSITE AND RUN UNCHANGED IN MATLAB 2013A WITH SOME METHODS' C++ MEX IMPLEMENTATION.

Method	Alg.1	Alg.2	LPS	UFO[1]	SVO[21]	CB[37]	PD[47]	HPS [13]	LR[16]	CA[8]	DS[45]	PD[47]	HPS[13]
Time(s)	0.87	9.56	2.45	18.73	40.33	1.18	3.64	5.02	11.92	36.05	0.84	19.45	3.16

TABLE II
PARAMETER SELECTION AND MODEL ROBUSTNESS. WE TEST DIFFERENT PARAMETERS ON THREE BENCHMARKS IN TERMS OF F-MEASURE (HIGHER IS BETTER) AND MAE (LOWER IS BETTER). THE BEST PARAMETERS ARE WRITTEN IN **BOLD**, WHICH ARE OUR MODEL'S DEFAULT SETTINGS. **RED** AND **BLUE** NUMBERS IN BOLD REPRESENT THE BEST AND BETTER PERFORMANCE IN EACH EVALUATION CATEGORY.

Dataset	Metric	# of switching labels p_1, p_2			The objectness criterion γ_1					The compactness criterion γ_2			
		2, 150	5, 175	10, 200	0.4	0.6	**0.8**	1.0	1.2	1.4	**1.6**	1.8	2.0
MSRA-1000	F-measure	**0.90**	**0.83**	0.52	0.72	0.84	**0.90**	**0.87**	0.80	0.83	**0.90**	**0.87**	0.79
	MAE	**0.07**	**0.13**	0.31	0.20	0.11	**0.07**	**0.09**	0.13	0.12	**0.07**	**0.09**	0.16
CCSD-1000	F-measure	**0.68**	**0.56**	0.55	0.60	**0.683**	**0.681**	0.64	0.58	0.63	**0.68**	**0.66**	0.60
	MAE	**0.23**	**0.32**	0.35	0.33	**0.21**	**0.23**	0.28	0.35	0.27	**0.23**	**0.26**	0.32
MSRA-5000	F-measure	**0.81**	**0.77**	0.63	0.75	0.79	**0.81**	**0.83**	0.78	0.77	**0.813**	0.78	**0.811**
	MAE	**0.12**	**0.25**	0.30	0.20	0.16	**0.12**	**0.11**	0.16	0.25	**0.122**	0.15	**0.124**

and recall. At last, the lowest MAE error is accomplished by LPS on this dataset.

5) PASCAL-S and THU-10,000: Fig.8 shows the performance comparison with other algorithms on the THU-10,000 and PASCAL-S benchmarks, in terms of a continuous-map (PR-curve) and an adaptive-threshold evaluation. We achieve comparable performance with the best results reported so far. Specifically, the F-measure and precision are the highest as well as MAE the lowest on the THU dataset; on the PASCAL-S, LPS is less inferior than MR[23] and MK[24] in terms of F-measure and overlap value while we achieve the highest precision and a comparable MAE result with the best ones (HS[38], DS[45]). Note that the MAE of CA[8] is quite high on every dataset because the size of their saliency maps are much smaller than the original images[2].

6) Execution Time: Tab.I shows the average execution time of processing one image in the MSRA-1000 dataset. Experiments are conducted on an Intel Core i7-3770 machine, equipped with 3.40GHz dominant frequency and 32 GB RAM. Alg.2 takes much longer than Alg.1 because the calculation of the objectness measure [20] is time consuming. By introducing a selection scheme using the compactness criterion, the computational efficiency of LPS has increased 74%. In contrast, those methods that directly utilize the objectness measure for each single image (UFO[1], SVO[21]) have suffered from poor efficiency as well as inferior P-R curves.

Note that some methods such as CB[37] and DS[45] have faster efficiency than ours; we believe an effective parallelized acceleration using GPU implementation on the compactness calculation and the pixel-wise saliency coherence at a pixel basis can substantially improve the computational efficiency.

7) Parameter Selection and Model Robustness: Tab.II shows the quantitative results using different parameter combinations. We choose the best qualified parameters in terms of F-measure and MAE on the MSRA (1000 or 5000) and CCSD datasets. Our algorithm takes the least number of parameters in order to better generalise on different datasets.

C. Visual Comparison

Several natural images with complex background are shown through Fig.9 to Fig.11 for visual comparison of our method w.r.t. the most recent state-of-the-arts. From these examples, we can see that most saliency detectors can effectively handle cases with relatively simple background and homogenous objects, such as the third and fourth row from the bottom in Fig.9, the first three rows in Fig.11, etc.

However, our model can tackle even more complicated scenarios, for example: (a) *cluttered background*: row 1,2,4 in Fig.10, row 7, 12 in Fig.11; (b) *low contrast between objects and background*: row 4,11,13 in Fig.11; (c) *heterogeneous objects*, row 2,10,11 in Fig.9; (d) *multi-scale objects*, the first three rows in Fig.9. More examples can be found in these figures. Due to the simple inner propagation process, our algorithm can effectively separate background labels and assign high saliency values to the dissimilar superpixels, *i.e.*, the candidate salient objects. With the help of a foreground proposal scheme, *i.e.*, objectness, the inter propagation can redirect the selection of foreground labels and compensate the intermediate results from the inner stage, thus detecting more accurate salient objects even from low contrast foreground and cluttered background.

D. Limitation and Analysis

Examples in the last rows of Fig.9 to Fig.11 show failure cases where the proposed algorithm is unable to detect the

[2] Most saliency maps are of size 300×400 as well as the size of ground truth maps; due to the multi-scale implementation in [8], the larger dimension of their maps is fixed to be 250. For fair comparison, we resize their smaller results to the same size of the ground truth maps and compute the corresponding evaluation metrics.

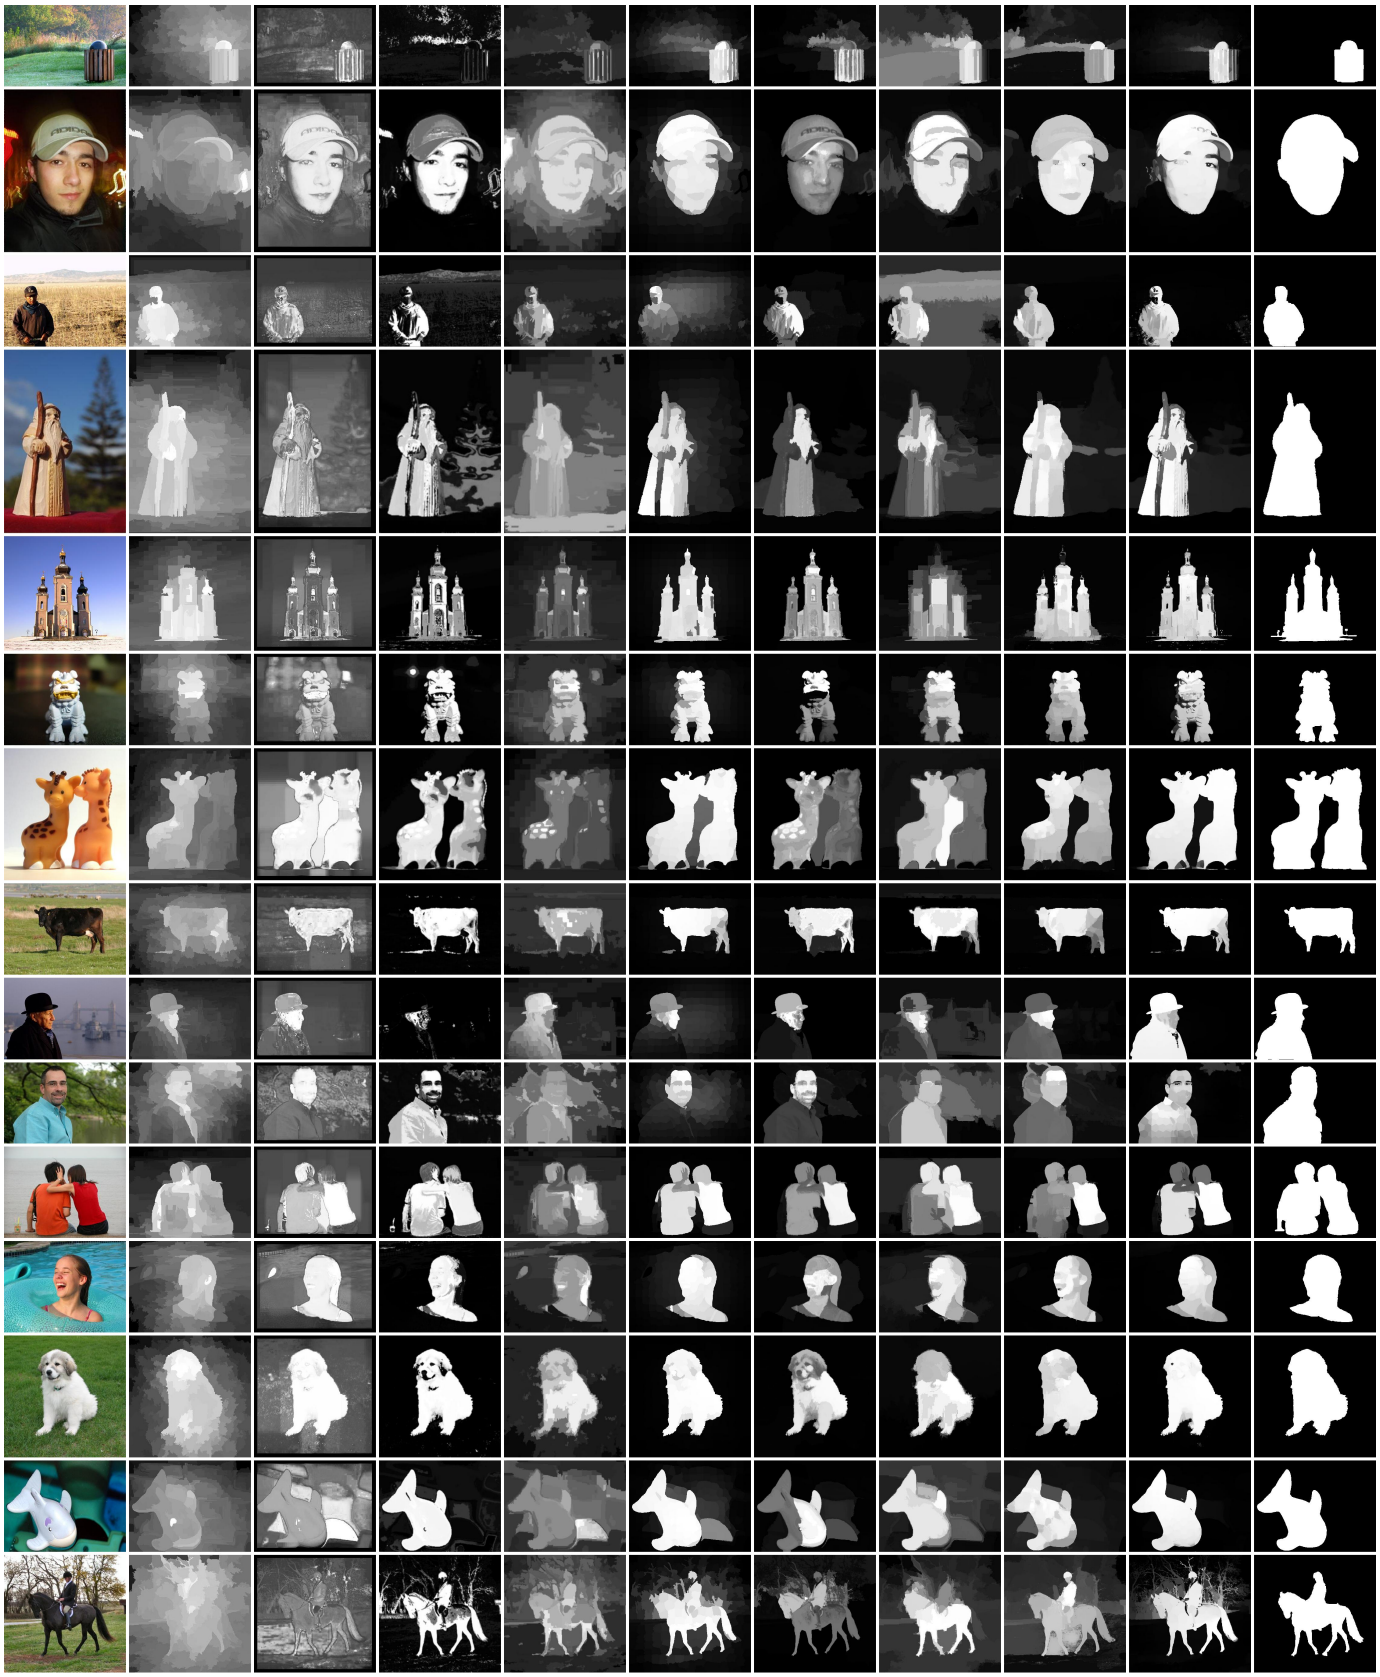

Fig. 9. Visual comparison of previous methods, our algorithm (LPS) and ground truth (GT) on the MSRA-5000 dataset. The last example shows a failure case where LPS overwhelmingly highlights the background around the horse due to a complex configuration of color and texture in the background.

Fig. 10. Visual comparison of previous methods, our algorithm (LPS) and ground truth (GT) on the CCSD-1000 dataset. The examples in the last two rows show failure cases where LPS abundantly detects more 'salient area' or powerlessly segments the salient object from the complex background.

salient object in some scenarios. Currently we only use the color information to construct the affinity matrix because the structure description of an image is included in the pre-abstraction processing. As shown in Sec.III-D, the structure based descriptor does not work well due to redundant extraction of the foreground and noisy extraction of the background. However, we believe that investigating more sophisticated feature representations for the co-transduction algorithm would be greatly beneficial. It would also be interesting to exploit top-down and category-independent semantic information to enhance the current results. We will leave these two directions as the starting point of our future research.

V. Conclusions

In this work, we explicitly propose a label propagation method in salient object detection. For some images, an inner label propagation via boundary labels alone obtains good visual and evaluation results; for more natural and complex images in the wild, a co-transduction algorithm which combines boundary superpixels with objectness labels can have better saliency assignment. The compactness criterion decides whether the final saliency map is simply a production of the inner propagation or a fusion outcome of the inter propagation. The proposed method achieves superior performance in terms of different evaluation metrics, compared with the state-of-the-arts on five benchmark image datasets.

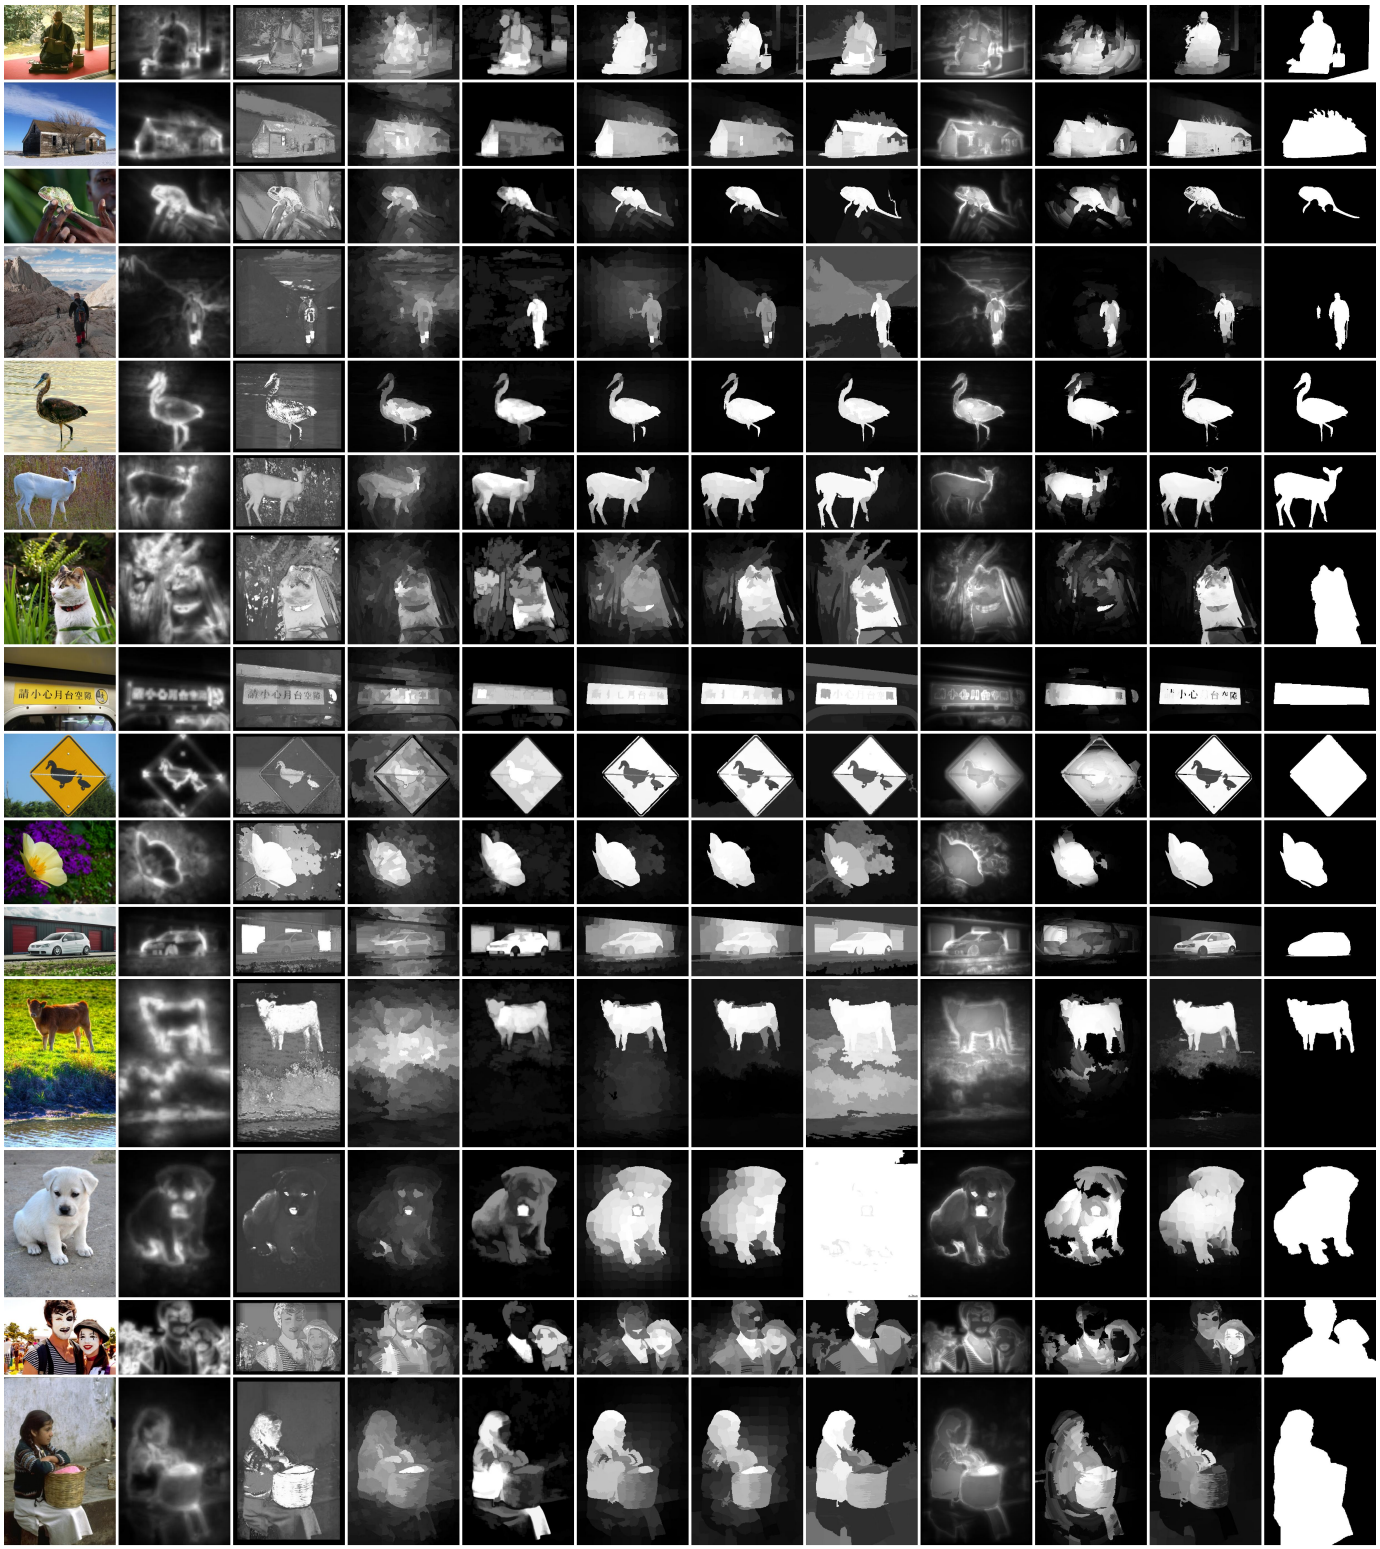

Fig. 11. Visual comparison of previous methods, our algorithm (LPS) and ground truth (GT) on the PASCAL-S and THUS-10,000 dataset. The examples in the last two rows show failure cases where LPS carelessly misses the foreground parts that belong to the salient people.

REFERENCES

[1] P. Jiang, H. Ling, J. Yu, and J. Peng, "Salient region detection by UFO: Uniqueness, Focusness and Objectness," in *ICCV*, 2013. 1, 5, 7, 8, 9, 10

[2] L. Itti, C. Koch, and E. Niebur, "A model of saliency-based visual attention for rapid scene analysis," *IEEE Trans. PAMI*, vol. 20, pp. 1254–1259, 1998. 1, 2, 7

[3] T. Judd, K. Ehinger, F. Durand, and A. Torralba, "Learning to predict where humans look," in *ICCV*, 2009. 1

[4] J. Zhang and S. Sclaroff, "Saliency detection: A boolean map approach." in *ICCV*, 2013. 1, 7, 8, 11, 12

[5] T. Liu, Z. Yuan, J. Sun, J. Wang, N. Zheng, X. Tang, and H. Shum, "Learning to detect a salient object," *IEEE Trans. PAMI*, vol. 33, 2011. 1, 4, 6

[6] R. Achanta, F. Estrada, P. Wils, and S. Süsstrunk, "Salient region detection and segmentation," *Computer Vision Systems*, pp. 66–75, 2008. 1

[7] V. Gopalakrishnan, Y. Hu, and D. Rajan, "Salient region detection by modeling distributions of color and orientation," *IEEE Trans. Multimedia*, vol. 11, no. 5, pp. 892–905, 2009. 1

[8] S. Goferman, L. Zelnik-Manor, and A. Tal, "Context-aware saliency detection," in *CVPR*, 2010. 1, 7, 9, 12

[9] L. Itti, "Automatic foveation for video compression using a neurobiological model of visual attention," *IEEE Trans. Image Proc.*, vol. 13, no. 10, pp. 1304–1318, Oct 2004. 1

[10] M.-M. Cheng, G.-X. Zhang, N. J. Mitra, X. Huang, and S.-M. Hu, "Global contrast based salient region detection," in *CVPR*, 2011. 1, 2, 3, 7, 10

[11] U. Rutishauser, D. Walther, C. Koch, and P. Perona, "Is bottom-up attention useful for object recognition?" in *CVPR*, 2004. 1

[12] V. Mahadevan and N. Vasconcelos, "Saliency-based discriminant tracking," in *CVPR*, 2009. 1

[13] X. Li, Y. Li, C. Shen, A. Dick, and A. Hengel, "Contextual hypergraph modeling for salient object detection," in *ICCV*, 2013. 1, 2, 7, 9, 11

[14] C. Scharfenberger, A. Wong, K. Fergani, J. S. Zelek, and D. A. Clausi, "Statistical textural distinctiveness for salient region detection in natural images," in *CVPR*, 2013. 1, 7, 8

[15] K. Shi, K. Wang, J. Lu, and L. Lin, "PISA: Pixelwise image saliency by aggregating complementary appearance contrast measures with spatial priors," in *CVPR*, 2013. 1, 5, 7, 8

[16] X. Shen and Y. Wu, "A unified approach to salient object detection via low rank matrix recovery," in *CVPR*, 2012. 1, 2, 7, 9, 12

[17] H. Jiang, J. Wang, Z. Yuan, Y. Wu, N. Zheng, and S. Li, "Salient object detection: A discriminative regional feature integration approach," in *CVPR*, 2013. 1, 2, 6

[18] Z. Jiang and L. S. Davis, "Submodular salient region detection," in *CVPR*, 2013. 1

[19] P. Siva, C. Russell, T. Xiang, and L. Agapito, "Looking beyond the image: Unsupervised learning for object saliency and detection," in *CVPR*, 2013. 1, 2, 7

[20] B. Alexe, T. Deselaers, and V. Ferrari, "Measuring the objectness of image windows," *IEEE Trans. PAMI*, vol. 34, no. 11, pp. 2189–2202, 2012. 1, 4, 8, 9

[21] K.-Y. Chang, T.-L. Liu, H.-T. Chen, and S.-H. Lai, "Fusing generic objectness and visual saliency for salient object detection," in *ICCV*, 2011. 1, 7, 9, 10

[22] Y. Jia and M. Han, "Category-independent object-level saliency detection," in *ICCV*, 2013. 1, 3

[23] C. Yang, L. Zhang, H. Lu, X. Ruan, and M.-H. Yang, "Saliency detection via graph-based manifold ranking," in *CVPR*, 2013. 1, 2, 3, 7, 9, 11, 12

[24] B. Jiang, L. Zhang, H. Lu, and M.-H. Yang, "Saliency detection via absorbing markov chain." in *ICCV*, 2013. 1, 2, 3, 7, 9, 10, 11, 12

[25] X. Bai, B. Wang, C. Yao, W. Liu, and Z. Tu, "Co-transduction for shape retrieval," *IEEE Trans. Image Proc.*, vol. 21, no. 5, pp. 2747–2757, 2012. 1, 5, 6

[26] Y.-F. Ma and H.-J. Zhang, "Contrast-based image attention analysis by using fuzzy growing," in *ACM Multimedia*, 2003. 2

[27] J. Harel, C. Koch, and P. Perona, "Graph-based visual saliency," in *NIPS*, 2007. 2, 7

[28] X. Hou and L. Zhang, "Saliency detection: A spectral residual approach," in *CVPR*, 2007. 2, 4, 7

[29] R. Achanta, S. Hemami, F. Estrada, and S. Süsstrunk, "Frequency-tuned salient region detection," in *CVPR*, 2009. 2, 6, 7, 8

[30] J. Li, M. D. Levine, X. An, and H. He, "Saliency detection based on frequency and spatial domain analysis," in *BMVC*, 2011. 2

[31] A. Borji, D. N. Sihite, and L. Itti, "Salient object detection: A benchmark," in *ECCV*, 2012. 2

[32] X. Bai, X. Yang, L. J. Latecki, W. Liu, and Z. Tu, "Learning context-sensitive shape similarity by graph transduction," *IEEE Trans. PAMI*, vol. 32, no. 5, pp. 861–874, 2010. 2, 3

[33] L. Xu, C. Lu, Y. Xu, and J. Jia, "Image smoothing via L_0 gradient minimization," *ACM Trans. Graph.*, vol. 30, no. 6, pp. 174:1–174:12, Dec. 2011. 2

[34] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Süsstrunk, "SLIC superpixels," EPFL, Tech. Rep. 149300, June 2010. 2

[35] X. Zhu, Z. Ghahramani, and J. Lafferty, "Semi-supervised learning using gaussian fields and harmonic functions," in *ICML*, 2003, pp. 912–919. 4

[36] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *Int. J. Comput. Vision*, vol. 59, no. 2, pp. 167–181, Sep. 2004. 4

[37] H. Jiang, J. Wang, Z. Yuan, T. Liu, N. Zheng, and S. Li, "Automatic salient object segmentation based on context and shape prior," in *BMVC*, 2011. 6, 7, 9, 10

[38] Q. Yan, L. Xu, J. Shi, and J. Jia, "Hierarchical saliency detection," in *CVPR*, 2013. 6, 7, 8, 9, 11, 12

[39] Y. Wei, F. Wen, W. Zhu, and J. Sun, "Geodesic saliency using background priors," in *ECCV*, 2012. 6, 7, 8

[40] Y. Li, X. Hou, C. Koch, J. M. Rehg, and A. L. Yuille, "The secrets of salient object segmentation," in *CVPR*, 2014. 6

[41] Y. Zhai and M. Shah, "Visual attention detection in video sequences using spatiotemporal cues," *ACM Multimedia*, pp. 815–824, 2006. 7

[42] E. Rahtu, J. Kannala, M. Salo, and J. Heikkilä, "Segmenting salient objects from images and videos," in *ECCV*, 2010. 7, 10, 12

[43] Y. Xie and H. Lu, "Visual saliency detection based on Bayesian model," in *ICIP*, 2011. 7

[44] F. Perazzi, P. Krähenbühl, Y. Pritch, and A. Hornung, "Saliency filters: Contrast based filtering for salient region detection," in *CVPR*, 2012. 7, 8

[45] X. Li, H. Lu, L. Zhang, X. Ruan, and M.-H. Yang, "Saliency detection via dense and sparse reconstruction," in *ICCV*, 2013. 7, 8, 9, 11, 12

[46] M.-M. Cheng, J. Warrell, W.-Y. Lin, S. Zheng, V. Vineet, and N. Crook, "Efficient salient region detection with soft image abstraction," in *ICCV*, 2013. 7, 8, 10

[47] R. Margolin, A. Tal, and L. Zelnik-Manor, "What makes a patch distinct?" in *CVPR*, 2013. 7, 8, 9, 11, 12

[48] Z. Liu, W. Zou, and O. Le Meur, "Saliency tree: A novel saliency detection framework," *IEEE Transactions on Image Processing*, vol. 23, no. 5, pp. 1937–1952, May 2014. 7, 8, 10

[49] Y. Fang, Z. Chen, W. Lin, and C.-W. Lin, "Saliency detection in the compressed domain for adaptive image retargeting," *IEEE Transactions on Image Processing*, vol. 21, no. 9, pp. 3888–3901, 2012. 7